arXiv:1612.08154v1 [cs.FL] 24 Dec 2016

FAMILIES OF DFAS AS ACCEPTORS OF ω-REGULAR LANGUAGES

DANA ANGLUIN, UDI BOKER, AND DANA FISMAN

Yale University, New Haven, CT, USA

Interdisciplinary Center (IDC), Herzliya, Israel

Ben-Gurion University, Beer-Sheva, Israel

Abstract. Families of DFAs (FDFAs) provide an alternative formalism for recognizing ω-regular languages. The motivation for introducing them was a desired correlation between the automaton states and right congruence relations, in a manner similar to the Myhill-Nerode theorem for regular languages. This correlation is beneficial for learning algorithms, and indeed it was recently shown that ω-regular languages can be learned from membership and equivalence queries, using FDFAs as the acceptors.

In this paper, we look into the question of how suitable FDFAs are for defining ω-regular languages. Specifically, we look into the complexity of performing Boolean operations, such as complementation and intersection, on FDFAs, the complexity of solving decision problems, such as emptiness and language containment, and the succinctness of FDFAs compared to standard deterministic and nondeterministic ω-automata.

We show that FDFAs enjoy the benefits of deterministic automata with respect to Boolean operations and decision problems. Namely, they can all be performed in nondeterministic logarithmic space. We provide polynomial translations of deterministic Büchi and co-Büchi automata to FDFAs and of FDFAs to nondeterministic Büchi automata (NBAs). We show that translation of an NBA to an FDFA may involve an exponential blowup. Last, we show that FDFAs are more succinct than deterministic parity automata (DPAs) in the sense that translating a DPA to an FDFA can always be done with only a polynomial increase, yet the other direction involves an inevitable exponential blowup in the worst case.

1. Introduction

The theory of finite-state automata processing infinite words was developed in the early sixties, starting with Büchi [Büc60] and Muller [Mul63], and motivated by problems in logic and switching theory. Today, automata for infinite words are extensively used in verification and synthesis of *reactive systems*, such as operating systems and communication protocols.

2012 ACM CCS: [**Theory of computation**]: Formal languages and automata theory—Automata over infinite objects.

Key words and phrases: Finite automata, Omega-Regular Languages.

* The present article extends [ABF16].

Research of the second author was supported by the Israel Science Foundation grant 1373/16.

Research of the third author was supported by US NSF grant CCF-1138996.

LOGICAL METHODS IN COMPUTER SCIENCE DOI:10.2168/LMCS-???

© D. Angluin, U. Boker, and D. Fisman Creative Commons

An automaton processing finite words makes its decision according to the last visited state. On infinite words, Büchi defined that a run is accepting if it visits a designated set of states infinitely often. Since then several other accepting conditions were defined, giving rise to various ω-automata, among which are Muller, Rabin, Streett and parity automata.

The theory of ω-regular languages is more involved than that of finite words. This was first evidenced by Büchi's observation that nondeterministic Büchi automata are more expressive than their deterministic counterpart. While for some types of ω-automata the nondeterministic and deterministic variants have the same expressive power, none of them possesses all the nice qualities of acceptors for finite words. In particular, none has a corresponding Myhill-Nerode theorem [Ner58], i.e. a direct correlation between the states of the automaton and the equivalence classes corresponding to the canonical right congruence of the recognized language.

The absence of a Myhill-Nerode like property in ω-automata has been a major drawback in obtaining learning algorithms for ω-regular languages, a question that has received much attention lately due to applications in verification and synthesis, such as black-box checking [PVY02], assume-guarantee reasoning [AvMN05, PGB+08, CCF+10, FKP11], error localization [WN05, CCK+15], regular model checking [VSVA05, NJ13], finding security bugs [RMSM09, CPPdR14, FJV14], programming networks [RMSM09, YAL14] and more. The reason is that learning algorithms typically build on this correspondence between the automaton and the right congruence.

Recently, several algorithms for learning an unknown ω-regular language were proposed, all using non-conventional acceptors. One uses a reduction due to [CNP94] named $L_\$$-automata of ω-regular languages to regular languages [FCTW08], and the others use a representation termed *families of* DFA*s* [AF14, LCZL16]. Both representations are founded on the following well known property of ω-regular languages: two ω-regular languages are equivalent iff they agree on the set of ultimately periodic words. An ultimately periodic word uv^ω, where $u \in \Sigma^*$ and $v \in \Sigma^+$, can be represented as a pair of finite words (u, v). Both $L_\$$-automata and families of DFAs process such pairs and interpret them as the corresponding ultimately periodic words. Families of DFAs have been shown to be up to exponentially more succinct than $L_\$$-automata [AF14].

A family of DFAs (FDFA) is composed of a *leading automaton* $\mathcal{Q}$ with no accepting states and for each state q of $\mathcal{Q}$, a *progress* DFA $\mathcal{P}_q$. Intuitively, the leading automaton is responsible for processing the non-periodic part u, and depending on the state q reached when $\mathcal{Q}$ terminated processing u, the respective progress DFA $\mathcal{P}_q$ processes the periodic part v, and determines whether the pair (u, v), which corresponds to uv^ω, is accepted. (The exact definition is more subtle and is provided in Section 3.) If the leading automaton has n states and the size of the maximal progress DFA is k, we say that the FDFA is of size (n, k). An earlier definition of FDFAs, given in [Kla94], provided a machine model for the *families of right congruences* of [MS97].[1] They were redefined in [AF14], where their acceptance criterion was adjusted, and their size was reduced by up to a quadratic factor. We follow the definition of [AF14].

[1]Another related formalism is of Wilke algebras [Wil91, Wil93], which are two-sorted algebras equipped with several operations. An ω-language over Σ^ω is ω-regular if and only if there exists a two-sorted morphism from Σ^∞ into a finite Wilke structure [Wil91]. A central difference between the FDFA theory and the algebraic theory of recognition by monoids, semigroups, ω-semigroups, and Wilke structures is that the former relates to right-congruences, while the latter is based on two-sided congruences.

In order for an FDFA to properly characterize an ω-regular language, it must satisfy the *saturation* property: considering two pairs (u, v) and (u', v'), if $uv^\omega = u'v'^\omega$ then either both (u, v) and (u', v') are accepted or both are rejected (cf. [CNP94, Saë90]). Saturated FDFAs are shown to exactly characterize the set of ω-regular languages. Saturation is a semantic property, and the check of whether a given FDFA is saturated is shown to be in PSPACE. Luckily, the FDFAs that result from the learning algorithm of [AF14] are guaranteed to be saturated.

Saturated FDFAs bring an interesting potential – they have a Myhill-Nerode like property, and while they are "mostly" deterministic, a nondeterministic aspect is hidden in the separation of the prefix and period parts of an ultimately periodic infinite word. This gives rise to the natural questions of how "dominant" are the determinism and nondeterminism in FDFAs, and how "good" are they for representing ω-regular languages. These abstract questions translate to concrete questions that concern the succinctness of FDFAs and the complexity of solving their decision problems, as these measures play a key role in the usefulness of applications built on top of them.

Our purpose in this paper is to analyze the FDFA formalism and answer these questions. Specifically, we ask: What is the complexity of performing the Boolean operations of complementation, union, and intersection on saturated FDFAs? What is the complexity of solving the decision problems of membership, emptiness, universality, equality, and language containment for saturated FDFAs? How succinct are saturated FDFAs, compared to deterministic and nondeterministic ω-automata?

We show that saturated FDFAs enjoy the benefits of deterministic automata with respect to Boolean operations and decision functions. Namely, the Boolean operations can be performed in logarithmic space, and the decision problems can be solved in nondeterministic logarithmic space. The constructions and algorithms that we use extend their counterparts on standard DFAs. In particular, complementation of saturated FDFAs can be obtained on the same structure, and union and intersection is done on a product of the two given structures. The correctness proof of the latter is a bit subtle.

As for the succinctness, which turns out to be more involved, we show that saturated FDFAs properly lie in between deterministic and nondeterministic ω-automata. We provide polynomial translations from deterministic ω-automata to FDFAs and from FDFAs to nondeterministic ω-automata, and show that an exponential state blowup in the opposite directions is inevitable in the worst case.

Specifically, a saturated FDFA of size (n, k) can always be transformed into an equivalent nondeterministic Büchi automaton (NBA) with $O(n^2k^3)$ states. As for the other direction, transforming an NBA with n states to an equivalent FDFA is shown to be in $2^{\Theta(n \log n)}$. This is not surprising since, as shown by Michel [Mic88], complementing an NBA involves a $2^{\Omega(n \log n)}$ state blowup, while FDFA complementation requires no state blowup.

Considering deterministic ω-automata, a Büchi or co-Büchi automaton (DBA or DCA) with n states can be transformed into an equivalent FDFA of size $(n, 2n)$, and a deterministic parity automaton (DPA) with n states and k colors can be transformed into an equivalent FDFA of size (n, kn). As for the other direction, since DBA and DCA do not recognize all the ω-regular languages, while saturated FDFAs do, a transformation from an FDFA to a DBA or DCA need not exist. Comparing FDFAs to DPAs, which do recognize all ω-regular languages, we get that FDFAs can be exponentially more succinct: We show a family of languages $\{L_n\}_{n\geq 1}$, such that for every n, there exists an FDFA of size $(n+1, n^2)$ for L_n, but any DPA

recognizing L_n must have at least 2^{n-1} states. (A deterministic Rabin or Streett automaton for L_n is also shown to be exponential in n, requiring at least $2^{\frac{n}{2}}$ states.)

2. Preliminaries

An *alphabet* Σ is a finite set of symbols. The set of finite words over Σ is denoted by Σ^*, and the set of infinite words, termed ω-words, over Σ is denoted by Σ^ω. As usual, we use x^*, x^+, and x^ω to denote finite, non-empty finite, and infinite concatenations of x, respectively, where x can be a symbol, a finite word, or a langugae. We use ϵ for the empty word and Σ^+ for $\Sigma^* \setminus \{\epsilon\}$. An infinite word w is *ultimately periodic* if there are two finite words $u \in \Sigma^*$ and $v \in \Sigma^+$, such that $w = uv^\omega$. A *language* is a set of finite words, that is, a subset of Σ^*, while an ω-language is a set of ω-words, that is, a subset of Σ^ω. For natural numbers i and j and a word w, we use $[i..j]$ for the set $\{i, i+1, \ldots, j\}$, $w[i]$ for the i-th letter of w, and $w[i..j]$ for the subword of w starting at the i-th letter and ending at the j-th letter, inclusive.

An *automaton* is a tuple $\mathcal{A} = \langle \Sigma, Q, \iota, \delta \rangle$ consisting of an alphabet Σ, a finite set Q of states, an initial state $\iota \in Q$, and a transition function $\delta : Q \times \Sigma \to 2^Q$. A run of an automaton on a finite word $v = a_1 a_2 \ldots a_n$ is a sequence of states $q_0, q_1, \ldots, q_n$ such that $q_0 = \iota$, and for each $i \geq 0$, $q_{i+1} \in \delta(q_i, a_i)$. A run on an infinite word is defined similarly and results in an infinite sequence of states. The transition function is naturally extended to a function $\delta : Q \times \Sigma^* \to 2^Q$, by defining $\delta(q, \epsilon) = \{q\}$, and $\delta(q, av) = \cup_{p \in \delta(q,a)} \delta(p, v)$ for $q \in Q$, $a \in \Sigma$, and $v \in \Sigma^*$. We often use $\mathcal{A}(v)$ as a shorthand for $\delta(\iota, v)$ and $|\mathcal{A}|$ for the number of states in Q. We use $\mathcal{A}^q$ to denote the automaton $\langle \Sigma, Q, q, \delta \rangle$ obtained from $\mathcal{A}$ by replacing the initial state with q. We say that $\mathcal{A}$ is *deterministic* if $|\delta(q, a)| \leq 1$ and *complete* if $|\delta(q, a)| \geq 1$, for every $q \in Q$ and $a \in \Sigma$. For simplicity, we consider all automata to be complete. (As is known, every automaton can be linearly translated to an equivalent complete automaton.)

By augmenting an automaton with an acceptance condition α, thereby obtaining a tuple $\langle \Sigma, Q, \iota, \delta, \alpha \rangle$, we get an *acceptor*, a machine that accepts some words and rejects others. An acceptor accepts a word if at least one of the runs on that word is accepting. For finite words the acceptance condition is a set $F \subseteq Q$ of *accepting states*, and a run on a word v is accepting if it ends in an accepting state, i.e., if $\delta(\iota, v)$ contains an element of F. For infinite words, there are various acceptance conditions in the literature; here we mention three: Büchi, co-Büchi, and parity.[2] The Büchi and co-Büchi acceptance conditions are also a set $F \subseteq Q$. A run of a Büchi automaton is accepting if it visits F infinitely often. A run of a co-Büchi automaton is accepting if it visits F only finitely many times. A parity acceptance condition is a map $\kappa : Q \to [1..k]$ assigning each state a color (or rank). A run is accepting if the minimal color visited infinitely often is odd. We use $[\![\mathcal{A}]\!]$ to denote the set of words accepted by a given acceptor $\mathcal{A}$, and say that $\mathcal{A}$ *accepts* or *recognizes* $[\![\mathcal{A}]\!]$. Two acceptors $\mathcal{A}$ and $\mathcal{B}$ are *equivalent* if $[\![\mathcal{A}]\!] = [\![\mathcal{B}]\!]$.

We use three letter acronyms to describe acceptors, where the first letter is either D or N depending on whether the automaton is *deterministic* or *nondeterministic*, respectively. The second letter is one of {F,B,C,P}: F if this is an acceptor over finite words, B, C, or P

[2]There are other acceptance conditions in the literature, the most known of which are weak, Rabin, Streett, and Muller. The three conditions that we concentrate on are the most used ones; Büchi and co-Büchi due to their simplicity, and parity due to being the simplest for which the deterministic variant is strong enough to express all ω-regular languages.

if it is an acceptor over infinite words with Büchi, co-Büchi, or parity acceptance condition, respectively. The third letter is always A for acceptor.

For finite words, NFA and DFA have the same expressive power. A language is said to be *regular* if it is accepted by an NFA. For infinite words, the theory is more involved. While NPAs, DPAs, and NBAs have the same expressive power, DBAs, NCAs, and DCAs are strictly weaker than NBAs. An ω-language is said to be *ω-regular* if it is accepted by an NBA.

3. Families of DFAs (FDFAs)

It is well known that two ω-regular languages are equivalent if they agree on the set of ultimately periodic words (this is a consequence of McNaughton's theorem [McN66]). An ultimately periodic word uv^ω, where $u \in \Sigma^*$ and $v \in \Sigma^+$, is usually represented by the pair (u, v). A canonical representation of an ω-regular language can thus consider only ultimately periodic words, namely define a language of pairs $(u, v) \in \Sigma^* \times \Sigma^+$. Such a representation $\mathcal{F}$ should satisfy the *saturation* property: considering two pairs (u, v) and (u', v'), if $uv^\omega = u'v'^\omega$ then either both (u, v) and (u', v') are accepted by $\mathcal{F}$ or both are rejected by $\mathcal{F}$.

A family of DFAs (FDFA) accepts such pairs (u, v) of finite words. Intuitively, it consists of a *leading automaton* $\mathcal{Q}$ with no acceptance condition that runs on the prefix-word u, and for each state q of $\mathcal{Q}$, a *progress automaton* $\mathcal{P}_q$, which is a DFA that runs on the period-word v.

A straightforward definition of acceptance for a pair (u, v), could have been that the run of the leading automaton $\mathcal{Q}$ on u ends at some state q, and the run of the progress automaton $\mathcal{P}_q$ on v is accepting. This goes along the lines of $L_\$$-automata [CNP94]. However, such an acceptance definition does not fit well the saturation requirement, and might enforce very large automata [AF14]. The intuitive reason is that every progress automaton might need to handle the period-words of all prefix-words.

To better fit the saturation requirement, the acceptance condition of an FDFA is defined with respect to a *normalization* of the input pair (u, v). The normalization is a new pair (x, y), such that $xy^\omega = uv^\omega$, and in addition, the run of the leading automaton $\mathcal{Q}$ on xy^i ends at the same state for every natural number i. Over the normalized pair (x, y), the acceptance condition follows the straightforward approach discussed above. This normalization resembles the implicit flexibility in the acceptance conditions of ω-automata, such as the Büchi condition, and allows saturated FDFAs to be up to exponentially more succinct than $L_\$$-automata [AF14].

Below, we formally define an FDFA, the normalization of an input pair (u, v), and the acceptance condition. We shall use Σ^{*+} as a shorthand for $\Sigma^* \times \Sigma^+$, whereby the input to an FDFA is a pair $(u, v) \in \Sigma^{*+}$.

Definition 3.1 (A family of DFAs (FDFA))**.** [3]

[3]The FDFAs defined here follow the definition in [AF14], which is a little different from the definition of FDFAs in [Kla94]; the latter provide a machine model for the families of right congruences introduced in [MS97]. The main differences between the two definitions are: i) In [Kla94], a pair (u, v) is accepted by an FDFA $\mathcal{F} = (\mathcal{Q}, \mathbf{P})$ if there is some factorization (x, y) of (u, v), such that $\mathcal{Q}(u) = q$ and $\mathcal{P}_q$ accepts v; and ii) in [Kla94], the FDFA $\mathcal{F}$ should also satisfy the constraint that for all words $u \in \Sigma^*$ and $v, v' \in \Sigma^+$, if $\mathcal{P}_{\mathcal{Q}(u)}(v) = \mathcal{P}_{\mathcal{Q}(u)}(v')$ then $\mathcal{Q}(uv) = \mathcal{Q}(uv')$.

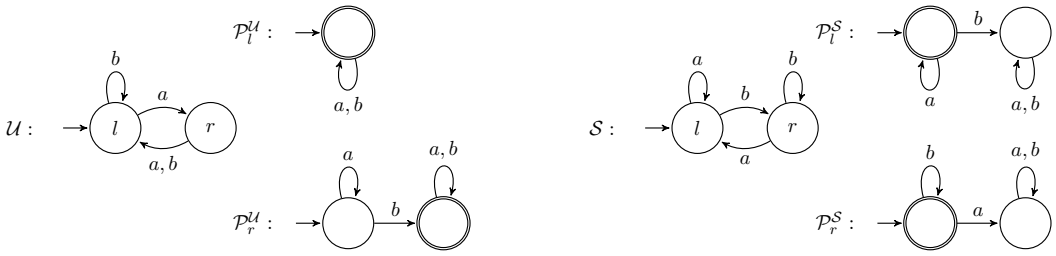

Figure 1: Left: an unsaturated FDFA with the leading automaton $\mathcal{U}$ and progress DFAs $\mathcal{P}_l^{\mathcal{U}}$ and $\mathcal{P}_r^{\mathcal{U}}$. Right: a saturated FDFA with the leading automaton $\mathcal{S}$ and progress DFAs $\mathcal{P}_l^{\mathcal{S}}$ and $\mathcal{P}_r^{\mathcal{S}}$.

- A *family of* DFA*s* (FDFA) is a pair $(\mathcal{Q}, \mathbf{P})$, where $\mathcal{Q} = (\Sigma, Q, \iota, \delta)$ is a deterministic *leading* automaton, and $\mathbf{P}$ is a set of $|\mathcal{Q}|$ DFAs, including for each state $q \in Q$, a *progress* DFA $\mathcal{P}_q = (\Sigma, P_q, \iota_q, \delta_q, F_q)$.
- Given a pair $(u, v) \in \Sigma^{*+}$ and an automaton $\mathcal{A}$, the *normalization* of (u, v) w.r.t $\mathcal{A}$ is the pair $(x, y) \in \Sigma^{*+}$, such that $x = uv^i$, $y = v^j$, and $i \geq 0$, $j \geq 1$ are the smallest numbers for which $\mathcal{A}(uv^i) = \mathcal{A}(uv^{i+j})$. (Since we consider complete automata, such a unique pair (x, y) is guaranteed.)
- Let $\mathcal{F} = (\mathcal{Q}, \mathbf{P})$ be an FDFA, $(u, v) \in \Sigma^{*+}$, and $(x, y) \in \Sigma^{*+}$ the normalization of (u, v) w.r.t $\mathcal{Q}$. We say that (u, v) is *accepted* by $\mathcal{F}$ iff $\mathcal{Q}(x) = q$ for some state q of $\mathcal{Q}$ and $\mathcal{P}_q(y)$ is an accepting state of $\mathcal{P}_q$.
- We use $[\![\mathcal{F}]\!]$ to denote the set of pairs accepted by $\mathcal{F}$.
- We define the *size* of $\mathcal{F}$, denoted by $|\mathcal{F}|$, as the pair $(|\mathcal{Q}|, \max\{|\mathcal{P}_q|\}_{q \in \mathcal{Q}})$, where $|\mathcal{A}|$ is the number of states in a DFA $\mathcal{A}$.
- An FDFA $\mathcal{F}$ is saturated if for every two pairs (u, v) and (u', v') such that $uv^\omega = u'v'^\omega$, either both (u, v) and (u', v') are in $[\![\mathcal{F}]\!]$ or both are not in $[\![\mathcal{F}]\!]$.

A saturated FDFA can be used to characterize an ω-regular language (see Theorem 5.2), while an unsaturated FDFA cannot.

An unsaturated FDFA is depicted in Figure 1 on the left. Consider the pairs (b, a) and (ba, aa). The former is normalized into (b, aa), as the run of the leading automaton $\mathcal{U}$ on "b" reaches the state l, and then when $\mathcal{U}$ iterates on "a", the first state to be revisited is l, which happens after two iterations. The latter is normalized into (ba, aa), namely it was already normalized. Now, (b, a) is accepted since in the run on its normalization (b, aa), $\mathcal{U}$ reaches the state l when running on "b", and the progress automaton $\mathcal{P}_l^{\mathcal{U}}$ accepts "aa". On the other hand, (ba, aa) is not accepted since the run of $\mathcal{U}$ on "ba" reaches the state r, and the progress automaton $\mathcal{P}_r^{\mathcal{U}}$ does not accept "aa". Yet, $ba^\omega = ba(aa)^\omega$, so the FDFA should have either accepted or rejected both (b, a) and (ba, aa), were it saturated, which is not the case.

A saturated FDFA is depicted in Figure 1 on the right. It accepts pairs of the forms (Σ^*, a^+) and (Σ^*, b^+), and characterizes the ω-regular language $(a + b)^*(a^\omega + b^\omega)$ of words in which there are eventually only a's or only b's.

4. Boolean Operations and Decision Procedures

We provide below algorithms for performing the Boolean operations of complementation, union, and intersection on saturated FDFAs, and deciding the basic questions on them, such as emptiness, universality, and language containment. All of these algorithms can be done in nondeterministic logarithmic space, taking advantage of the partial deterministic nature of FDFAs.[4] We conclude the section with the decision problem of whether an arbitrary FDFA is saturated, showing that it can be resolved in polynomial space.

Boolean operations. Saturated FDFAs are closed under Boolean operations as a consequence of Theorem 5.2, which shows that they characterize exactly the set of ω-regular languages. We provide below explicit algorithms for these operations, showing that they can be done effectively.

Complementation of an FDFA is simply done by switching between accepting and non-accepting states in the progress automata, as is done with DFAs.

Theorem 4.1. *Let $\mathcal{F}$ be an* FDFA*. There is a constant-space algorithm to obtain an* FDFA *$\mathcal{F}^c$, such that $[\![\mathcal{F}^c]\!] = \Sigma^{*+} \setminus [\![\mathcal{F}]\!]$, $|\mathcal{F}^c| = |\mathcal{F}|$, and $\mathcal{F}^c$ is saturated iff $\mathcal{F}$ is.*

Proof. Let $\mathcal{F} = (\mathcal{Q}, \mathbf{P})$, where for each state q of $\mathcal{Q}$, $\mathbf{P}$ has the DFA $\mathcal{P}_q = (\Sigma, P_q, \iota_q, \delta_q, F_q)$. We define $\mathcal{F}^c$ to be the FDFA $(\mathcal{Q}, \mathbf{P}^c)$, where for each state q of $\mathcal{Q}$, $\mathbf{P}^c$ has the DFA $\mathcal{P}^c_q = (\Sigma, P_q, \iota_q, \delta_q, P_q \setminus F_q)$. We claim that $\mathcal{F}^c$ recognizes the complement language of $\mathcal{F}$. Indeed, let $(u, v) \in \Sigma^{*+}$ and (x, y) its normalization with respect to $\mathcal{Q}$. Then $(u, v) \in [\![\mathcal{F}]\!]$ iff $y \in [\![\mathcal{P}_{\mathcal{Q}(x)}]\!]$. Thus $(u, v) \notin [\![\mathcal{F}]\!]$ iff $y \notin [\![\mathcal{P}_{\mathcal{Q}(x)}]\!]$ iff $y \in [\![\mathcal{P}^c_{\mathcal{Q}(x)}]\!]$ iff $(u, v) \in [\![\mathcal{F}^c]\!]$.

Since $\mathcal{F}$ is saturated, so is $\mathcal{F}^c$, as for all pairs (u, v) and (u', v') such that $uv^\omega = u'v'^\omega$, $\mathcal{F}$ either accepts or rejects them both, implying that $\mathcal{F}^c$ either rejects or accepts them both, respectively. □

Union and intersection of saturated FDFAs also resemble the case of DFAs, and are done by taking the product of the leading automata and each pair of progress automata. Yet, the correctness proof is a bit subtle, and relies on the following lemma, which shows that for a normalized pair (x, y), the period-word y can be manipulated in a certain way, while retaining normalization.

Lemma 4.2. *Let $\mathcal{Q}$ be an automaton, and let (x, y) be the normalization of some $(u, v) \in \Sigma^{*+}$ w.r.t. $\mathcal{Q}$. Then for every $i \geq 0$, $j \geq 1$ and finite words y', y'' such that $y = y'y''$, we have that $(xy^iy', (y''y')^j)$ is the normalization of itself w.r.t. Q.*

Proof. Let $x_1 = xy^iy'$ and $y_1 = (y''y')^j$. Since (x, y) is normalized w.r.t. $\mathcal{Q}$, we know that the run of $\mathcal{Q}$ on xy^ω is of the form $q_0, q_1, \ldots, q_{k-1}, (q_k, q_{k+1}, q_{k+2}, \ldots, q_m)^\omega$, where $|x| = k$ and $|y| = (m - k) + 1$. As $xy^\omega = x_1y_1^\omega$, the run of $\mathcal{Q}$ on $x_1y_1^\omega$ is identical. Since x is a prefix of x_1, the position $|x_1|$ lies within the repeated period, implying that $|x_1|$ is the first position, from $|x_1|$ onwards, that is repeated along the aforementioned run. Since y_1 is a cyclic repetition of y, and $\mathcal{Q}$ loops back over y, it also loops back over y_1. Thus (x_1, y_1) is the normalization of itself w.r.t. Q. □

[4] Another model that lies in between deterministic and nondeterministic automata are "semi-deterministic Büchi automata" [VW86], which are Büchi automata that are deterministic in the limit: from every accepting state onward, their behaviour is deterministic. Yet, as opposed to FDFAs, complementation of semi-deterministic Büchi automata might involve an exponential state blowup [BHS+16].

We continue with the union and intersection of saturated FDFAs.

Theorem 4.3. *Let $\mathcal{F}_1$ and $\mathcal{F}_2$ be saturated* FDFA*s of size (n_1, k_1) and (n_2, k_2), respectively. There exist logarithmic-space algorithms to obtain saturated* FDFA*s $\mathcal{H}$ and $\mathcal{H}'$ of size $(n_1 n_2,\ k_1 k_2)$, such that $[\![\mathcal{H}]\!] = [\![\mathcal{F}_1]\!] \cap [\![\mathcal{F}_2]\!]$ and $[\![\mathcal{H}']\!] = [\![\mathcal{F}_1]\!] \cup [\![\mathcal{F}_2]\!]$.*

Proof. The constructions of the union and intersection FDFAs are similar, only differing by the accepting states. We shall thus describe them together.

Construction. Given two automata $\mathcal{A}_1$ and $\mathcal{A}_2$, where $\mathcal{A}_i = (\Sigma, A_i, \iota_i, \delta_i)$, we denote by $\mathcal{A}_1 \times \mathcal{A}_2$ the product automaton $(\Sigma, A_1 \times A_2, (\iota_1, \iota_2), \delta_\times)$, where for every $\sigma \in \Sigma$, $\delta_\times((q_1, q_2), \sigma) = (\delta(q_1, \sigma), \delta(q_2, \sigma))$.

Given two DFAs $\mathcal{D}_1 = (\mathcal{A}_1, F_1)$ and $\mathcal{D}_2 = (\mathcal{A}_2, F_2)$, over the automata $\mathcal{A}_1$ and $\mathcal{A}_2$, and with the accepting states F_1 and F_2, respectively, we define the DFAs $\mathcal{D}_1 \otimes \mathcal{D}_2$ and $\mathcal{D}_1 \oplus \mathcal{D}_2$ as follows:

- $\mathcal{D}_1 \otimes \mathcal{D}_2 = (\mathcal{A}_1 \times \mathcal{A}_2, F_1 \times F_2)$
- $\mathcal{D}_1 \oplus \mathcal{D}_2 = (\mathcal{A}_1 \times \mathcal{A}_2, F_1 \times A_2 \cup A_1 \times F_2)$

Given two sets of DFAs $\mathbf{P}_1$ and $\mathbf{P}_2$, we define the sets of DFAs $\mathbf{P}_1 \otimes \mathbf{P}_2$ and $\mathbf{P}_1 \oplus \mathbf{P}_2$ as follows:

- $\mathbf{P}_1 \otimes \mathbf{P}_2 = \{\mathcal{D}_1 \otimes \mathcal{D}_2 \mid \mathcal{D}_1 \in \mathbf{P}_1 \text{ and } \mathcal{D}_2 \in \mathbf{P}_2\}$
- $\mathbf{P}_1 \oplus \mathbf{P}_2 = \{\mathcal{D}_1 \oplus \mathcal{D}_2 \mid \mathcal{D}_1 \in \mathbf{P}_1 \text{ and } \mathcal{D}_2 \in \mathbf{P}_2\}$

Given saturated FDFAs $\mathcal{F}_1 = (\mathcal{Q}_1, \mathbf{P}_1)$ and $\mathcal{F}_2 = (\mathcal{Q}_2, \mathbf{P}_2)$, we claim that $\mathcal{H} = (\mathcal{Q}_1 \times \mathcal{Q}_2, \mathbf{P}_1 \otimes \mathbf{P}_2)$ and $\mathcal{H}' = (\mathcal{Q}_1 \times \mathcal{Q}_2, \mathbf{P}_1 \oplus \mathbf{P}_2)$ are saturated FDFAs that recognize the intersection and union of $[\![\mathcal{F}_1]\!]$ and $[\![\mathcal{F}_2]\!]$, respectively.

Notice that the number of states in $\mathcal{H}$ and $\mathcal{H}'$ is quadratic in the number of states in $\mathcal{F}_1$ and $\mathcal{F}_2$, yet their representations is very similar to the representations of $\mathcal{F}_1$ and $\mathcal{F}_2$, for which reason the construction can be done in logarithmic space.

Correctness. Consider a pair $(u, v) \in \Sigma^{*+}$. Let (x_1, y_1) and (x_2, y_2) be its normalization with respect to $\mathcal{Q}_1$ and $\mathcal{Q}_2$, respectively, where $x_1 = uv^{i_1}$, $y_1 = v^{j_1}$, $x_2 = uv^{i_2}$, and $y_2 = v^{j_2}$. Let $i = \max(i_1, i_2)$ and j be the least common multiple of (j_1, j_2). Define $x = uv^i$ and $y = v^j$. One can verify that the normalization of (u, v) with respect to $\mathcal{Q}_1 \times \mathcal{Q}_2$ is (x, y).

We have $\mathcal{Q}_1 \times \mathcal{Q}_2\,(xy) = \mathcal{Q}_1 \times \mathcal{Q}_2\,(x) = (\mathcal{Q}_1(x), \mathcal{Q}_2(x))$. Since $xy^\omega = x_1 y_1^\omega$ and $\mathcal{F}_1$ is saturated, we get that $(x, y) \in [\![\mathcal{F}_1]\!]$ iff $(x_1, y_1) \in [\![\mathcal{F}_1]\!]$. Since the pair (x, y) satisfies the requirements of Lemma 4.2 w.r.t. (x_1, y_1) and $\mathcal{Q}_1$, it follows that (x, y) is a normalization of itself w.r.t. $\mathcal{Q}_1$. Thus, $y \in \mathcal{P}_{\mathcal{Q}_1(x)}$ iff $y_1 \in \mathcal{P}_{\mathcal{Q}_1(x_1)}$. Analogously, $y \in \mathcal{P}_{\mathcal{Q}_2(x)}$ iff $y_2 \in \mathcal{P}_{\mathcal{Q}_2(x_2)}$.

Hence, $(u, v) \in [\![\mathcal{F}_1]\!] \cap [\![\mathcal{F}_2]\!]$ iff $(y_1 \in \mathcal{P}_{\mathcal{Q}_1(x_1)}$ and $y_2 \in \mathcal{P}_{\mathcal{Q}_2(x_2)})$ iff $y \in \mathcal{P}_{\mathcal{Q}_1(x)} \otimes \mathcal{P}_{\mathcal{Q}_2(x)}$ iff $(u, v) \in \mathcal{H}$. Similarly, $(u, v) \in [\![\mathcal{F}_1]\!] \cup [\![\mathcal{F}_2]\!]$ iff $(y_1 \in \mathcal{P}_{\mathcal{Q}_1(x_1)}$ or $y_2 \in \mathcal{P}_{\mathcal{Q}_2(x_2)})$ iff $y \in \mathcal{P}_{\mathcal{Q}_1(x)} \oplus \mathcal{P}_{\mathcal{Q}_2(x)}$ iff $(u, v) \in \mathcal{H}'$.

The saturation of $\mathcal{H}$ and $\mathcal{H}'$ directly follows from the above proof of the languages they recognize: consider two pairs (u, v) and (u', v'), such that $uv^\omega = u'v'^\omega$. Then, by the saturation of $\mathcal{F}_1$ and $\mathcal{F}_2$, both pairs either belong, or not, to each of $[\![\mathcal{F}_1]\!]$ and $[\![\mathcal{F}_2]\!]$. Hence, both pairs belong, or not, to each of $[\![\mathcal{H}]\!] = [\![\mathcal{F}_1]\!] \cap [\![\mathcal{F}_2]\!]$ and $[\![\mathcal{H}']\!] = [\![\mathcal{F}_1]\!] \cup [\![\mathcal{F}_2]\!]$. □

Decision procedures. All of the basic decision problems can be resolved in nondeterministic logarithmic space, using the Boolean operations above and corresponding decision algorithms for DFAs.

The first decision question to consider is that of *membership*: given a pair (u, v) and an FDFA $\mathcal{F} = (\mathcal{Q}, \mathbf{P})$, does $\mathcal{F}$ accept (u, v)? The question is answered by normalizing (u, v) into a pair (x, y) and evaluating the runs of $\mathcal{Q}$ over x and of $\mathcal{P}_{\mathcal{Q}(x)}$ over y. A normalized pair is determined by traversing along $\mathcal{Q}$, making up to $|Q|$ repetitions of v. Notice that memory wise, x and y only require a logarithmic amount of space, as they are of the form $x = uv^i$ and $y = v^j$, where the representation of i and j is bounded by $\log |Q|$. The overall logarithmic-space solution follows from the complexity of algorithms for deterministically traversing along an automaton.

Proposition 4.4. *Given a pair $(u, v) \in \Sigma^{*+}$ and an FDFA $\mathcal{F}$ of size (n, k), the membership question, of whether $(u, v) \in [\![\mathcal{F}]\!]$, can be resolved in deterministic space of $O(\log n + \log k)$.*

The next questions to consider are those of *emptiness* and *universality*, namely given an FDFA $\mathcal{F} = (\mathcal{Q}, \mathbf{P})$, whether $[\![\mathcal{F}]\!] = \emptyset$, and whether $[\![\mathcal{F}]\!] = \Sigma^{*+}$, respectively. Notice that the universality problem is equivalent to the emptiness problem over the complement of $\mathcal{F}$. For nondeterministic automata, the complement automaton might be exponentially larger than the original one, making the universality problem much harder than the emptiness problem. Luckily, FDFA complementation is done in constant space, as is the case with deterministic automata, making the emptiness and universality problems equally easy.

The emptiness problem for an FDFA $(\mathcal{Q}, \mathbf{P})$ cannot be resolved by only checking whether there is a nonempty progress automaton in $\mathbf{P}$, since it might be that the accepted period v is not part of any normalized pair. Yet, the existence of a prefix-word x and a period-word y, such that $\mathcal{Q}(x) = \mathcal{Q}(xy)$ and $\mathcal{P}_{\mathcal{Q}(x)}$ accepts y is a sufficient and necessary criterion for the nonemptiness of $\mathcal{F}$. This can be tested in NLOGSPACE. Hardness in NLOGSPACE follows by a reduction from graph reachability [Jon75].

Theorem 4.5. *Emptiness and universality for FDFAs are NLOGSPACE-complete.*

Proof. An FDFA $\mathcal{F} = (\mathcal{Q}, \mathbf{P})$ is not empty iff there exists a pair $(u, v) \in \Sigma^{*+}$, whose normalization is some pair (x, y), such that $\mathcal{P}_{\mathcal{Q}(x)}$ accepts y. As a normalized pair is a normalization of itself, a sufficient and necessary criterion for the nonemptiness of $\mathcal{F}$ is the existence of a pair (x, y), such that $\mathcal{Q}(x) = \mathcal{Q}(xy)$ and $\mathcal{P}_{\mathcal{Q}(x)}$ accepts y.

We can nondeterministically find such a pair (x, y) in logarithmic space by guessing x and y (a single letter at each step), and traversing along $\mathcal{Q}$ and $\mathcal{P}_{\mathcal{Q}(x)}$ [GJ79].

Hardness in NLOGSPACE follows by a reduction from graph reachability [Jon75].

As FDFA complementation is done in constant space (Theorem 4.1), the universality problem has the same space complexity. □

The last decision questions we handle are those of *equality* and *containment*, namely given saturated FDFAs $\mathcal{F}$ and $\mathcal{F}'$, whether $[\![\mathcal{F}]\!] = [\![\mathcal{F}']\!]$ and whether $[\![\mathcal{F}]\!] \subseteq [\![\mathcal{F}']\!]$, respectively. Equality reduces to containment, as $[\![\mathcal{F}]\!] = [\![\mathcal{F}']\!]$ iff $[\![\mathcal{F}]\!] \subseteq [\![\mathcal{F}']\!]$ and $[\![\mathcal{F}']\!] \subseteq [\![\mathcal{F}]\!]$. Containment can be resolved by intersection, complementation, and emptiness check, as $[\![\mathcal{F}]\!] \subseteq [\![\mathcal{F}']\!]$ iff $[\![\mathcal{F}]\!] \cap [\![\mathcal{F}']\!]^c = \emptyset$. Hence, by Theorems 4.1, 4.3, and 4.5, these problems are NLOGSPACE-complete. Note that NLOGSPACE hardness immediately follows by reduction from the emptiness problem, which asks whether $[\![\mathcal{F}]\!] = \emptyset$. The complexity for *equality* and *containment* is easily derived from that of emptiness, intersection and complementation.

Proposition 4.6. *Equality and containment for saturated FDFAs are NLOGSPACE-complete.*

Saturation check. All of the operations and decision problems above assumed that the given FDFAs are saturated. This is indeed the case when learning FDFAs via the algorithm of [AF14], and when translating ω-automata to FDFAs (see Section 5). We show below that the decision problem of whether an arbitrary FDFA is saturated is in PSPACE. We leave the question of whether it is PSPACE-complete open.

Theorem 4.7. *The problem of deciding whether a given* FDFA *is saturated is in PSPACE.*

Proof. Let $\mathcal{F} = (\mathcal{Q}, \mathbf{P})$ be an FDFA of size (n, k). We first show that if $\mathcal{F}$ is unsaturated then there exist words u, v', v'' such that $|u| \leq n$ and $|v'|, |v''| \leq n^n k^{2k}$, and non-negative integers $l, r \leq k$ such that $(u, (v'v'')^l) \in \mathcal{F}$ while $(uv', (v''v')^r) \notin \mathcal{F}$.

If $\mathcal{F}$ is unsaturated then there exists some ultimately periodic word $w \in \Sigma^\omega$ that has two different decompositions to prefix and periodic words on which $\mathcal{F}$ provides different answers. Let $\mathcal{P}$ and $\mathcal{P}'$ be the respective progress automata, corresponding to states q and q' of $\mathcal{Q}$. Let the run of $\mathcal{Q}$ on w be $q_0, q_1, q_2, \ldots$. Since w is ultimately periodic, there exist $i, j \in \mathbb{N}$ such that $q_{h+j} = q_h$ for all $h > i$. That is, eventually the run cycles through a certain sequence of states. Then q and q' must be two states on the cycle where w settles. Let v' and v'' be the subwords of w that are read on the part of the shortest such cycle from q to q' and from q' back to q, respectively. Then the different decompositions are of the form $(u, (v'v'')^l)$ and $(uv', (v''v')^r)$ where u is a string that takes Q to q. Let l and r be the shortest such, then since $\mathcal{P}$ and $\mathcal{P}'$ have at most k states, we can assume $l, r \leq k$. We can also assume u is a shortest such string and thus $|u| \leq n$.

For a DFA $\mathcal{A} = \langle \Sigma, Q, \iota, \delta, F \rangle$ and a word $v \in \Sigma^*$, we use $\chi^{\mathcal{A}}_v$ to denote the function from Q to Q defined as $\chi^{\mathcal{A}}_v(q) = \delta(q, v)$. Note that given $|Q| = n$ there are at most n^n different such functions. Let $\mathcal{X} = \{\chi_v \mid \chi_v = \langle \chi^{\mathcal{Q}}_v, \chi^{\mathcal{P}}_v, \chi^{\mathcal{P}'}_v \rangle,\ v \in \Sigma^*\}$. Then $\mathcal{X}$ is the set of congruence classes of the relation $v_1 \approx_{\mathcal{X}} v_2$ iff $\chi^{\mathcal{Q}}_{v_1} = \chi^{\mathcal{Q}}_{v_2}$, $\chi^{\mathcal{P}}_{v_1} = \chi^{\mathcal{P}}_{v_2}$, and $\chi^{\mathcal{P}'}_{v_1} = \chi^{\mathcal{P}'}_{v_2}$. We can build an automaton such that each state corresponds to a class in $\mathcal{X}$, the initial state is χ_ϵ, and the transition relation is $\delta_{\mathcal{X}}(\chi_w, \sigma) = \chi_{w\sigma}$. The cardinality of $\mathcal{X}$ is at most $n^n k^{2k}$. Thus, every state has a representative word of length at most $n^n k^{2k}$ taking the initial state to that state. Therefore, there exist words y', y'' such that $y' \approx_{\mathcal{X}} v'$ and $y'' \approx_{\mathcal{X}} v''$ and $|y'|, |y''| \leq n^n k^{2k}$. Thus u, y', y'' and l, r satisfy the promised bounds.

Now, to see that FDFA saturation is in PSPACE note we can construct an algorithm that guesses integers $l, r \leq k$ and words u, v', v'' such that $|u| \leq n$ and $|v'|, |v''| \leq n^n k^{2k}$. It guesses the words letter by letter and constructs on the way $\chi_{v'v''}$ and $\chi_{v''v'}$. It also constructs along the way the states q' and q such that $q = \delta(u)$ and $q' = \delta(uv')$. It then computes the l and r powers of $\chi_{v'v''}$ and $\chi_{v'',v'}$, respectively. Finally, it checks whether one of $(\chi^{\mathcal{P}}_{v'v''})^l(q)$ and $(\chi^{\mathcal{P}'}_{v''v'})^r(q')$ is accepting and the other is not, and if so returns that $\mathcal{F}$ is unsaturated. The required space is $O(nk^2 \log nk^2)$. This shows that saturation is in coNPSPACE, and by Savitch's and Immerman-Szelepcsényi's theorems, in PSPACE. □

5. Translating To and From ω-Automata

As two ω-regular languages are equivalent iff they agree on the set of ultimately periodic words [McN66], an ω-regular language can be characterized by a language of pairs of finite words, and in particular by a saturated FDFA. We shall write $L \equiv L'$ to denote that a language $L \subseteq \Sigma^{*+}$ characterizes an ω-regular language L'. Formally:

Definition 5.1. A language $L \subseteq \Sigma^{*+}$ *characterizes* an ω-regular language $L' \subseteq \Sigma^\omega$, denoted by $L \equiv L'$, if for every pair $(u, v) \in L$, we have $uv^\omega \in L'$, and for every ultimately periodic word $uv^\omega \in L'$, we have $(u, v) \in L$.

The families of DFAs defined in [Kla94], as well as the analogous families of right congruences of [MS97], are known to characterize exactly the set of ω-regular languages [Kla94, MS97]. This is also the case with our definition of saturated FDFAs.

Theorem 5.2. *Every saturated* FDFA *characterizes an ω-regular language, and for every ω-regular language, there is a saturated* FDFA *characterizing it.*

Proof. The two directions are proved in Theorems 5.4 and 5.8, below. □

In this section, we analyze the state blowup involved in translating deterministic and nondeterministic ω-automata into equivalent saturated FDFAs, and vice versa. For nondeterministic automata, we consider the Büchi acceptance condition, since it is the simplest and most commonly used among all acceptance conditions. For deterministic automata, we consider the parity acceptance condition since it is the simplest among all acceptance conditions whose deterministic version is equi-expressible to the ω-regular languages. We also consider deterministic Büchi and co-Büchi, for the simple sub-classes they recognize.

5.1. **From ω-Automata to FDFAs.** We show that DBA, DCA, and DPA have polynomial translations to saturated FDFAs, whereas translation of NBAs to FDFAs may involve an inevitable exponential blowup.

From deterministic ω-automata. The constructions of a saturated FDFA that characterize a given DBA, DCA, or DPA $\mathcal{D}$ share the same idea: The leading automaton is equivalent to $\mathcal{D}$, except for ignoring the acceptance condition, and each progress automaton consists of several copies of $\mathcal{D}$, memorizing the acceptance level of the period-word. For a DBA or a DCA, two such copies are enough, memorizing whether or not a Büchi (co-Büchi) accepting state was visited. For a DPA with k colors, k such copies are required.

We start with the constructions of an FDFA for a given DBA or DCA, which are almost the same.

Theorem 5.3. *Let $\mathcal{D}$ be a* DBA *or a* DCA *with n states. There exists a saturated* FDFA $\mathcal{F}$ *of size $(n, 2n)$, such that $[\![\mathcal{F}]\!] \equiv [\![\mathcal{D}]\!]$.*

Proof.

Construction. Let $\mathcal{D} = \langle \Sigma, Q, \iota, \delta, \alpha \rangle$ be a DBA or a DCA. We define the FDFA $\mathcal{F} = (\mathcal{Q}, \mathbf{P})$, where $\mathcal{Q}$ is the same as $\mathcal{D}$ (without acceptance), and each progress automaton $\mathcal{P}_q$ has two copies of $\mathcal{D}$, having $(q, 0)$ as its initial state, and moving from the first to the second copy upon visiting a $\mathcal{D}$-accepting state. Formally: $\mathcal{Q} = \langle \Sigma, Q, \iota, \delta \rangle$, and for each state $q \in Q$, $\mathbf{P}$ has the DFA $\mathcal{P}_q = \langle \Sigma, Q \times \{0, 1\}, (q, 0), \delta', F \rangle$, where for every $\sigma \in \Sigma$, $\delta'((q, 0), \sigma) = (\delta(q, \sigma), 0)$ if $\delta(q, \sigma) \notin \alpha$ and $(\delta(q, \sigma), 1)$ otherwise; and $\delta'((q, 1), \sigma) = (\delta(q, \sigma), 1)$. The set F of accepting states is $Q \times \{1\}$ if $\mathcal{D}$ is a DBA and $Q \times \{0\}$ if $\mathcal{D}$ is a DCA.

Correctness. We show the correctness for the case that $\mathcal{D}$ is a DBA. The case that $\mathcal{D}$ is a DCA is analogous.

Consider a word $uv^\omega \in [\![\mathcal{D}]\!]$, and let (x, y) be the normalization of (u, v) w.r.t. $\mathcal{Q}$. Since $xy^\omega = uv^\omega \in [\![\mathcal{D}]\!]$, it follows that $\mathcal{D}$ visits an accepting state when running on y from the state $\mathcal{D}(x)$, implying that $\mathcal{P}_{\mathcal{Q}(x)}(y)$ is an accepting state. Hence, $(u, v) \in [\![\mathcal{F}]\!]$.

As for the other direction, consider a pair $(u, v) \in [\![\mathcal{F}]\!]$, and let (x, y) be the normalization of (u, v) w.r.t. $\mathcal{Q}$. Since $\mathcal{P}_{\mathcal{Q}(x)}(y)$ is an accepting state, it follows that $\mathcal{D}$ visits an accepting state when running on y from the state $\mathcal{D}(x)$, implying that $xy^\omega = uv^\omega \in [\![\mathcal{D}]\!]$.

Note that $\mathcal{F}$ is saturated as a direct consequence of the proof that it characterizes an ω-regular language. □

We continue with the construction of an FDFA for a given DPA.

Theorem 5.4. *Let $\mathcal{D}$ be a* DPA *with n states and k colors. There exists a saturated* FDFA *$\mathcal{F}$ of size (n, kn), such that $[\![\mathcal{F}]\!] \equiv [\![\mathcal{D}]\!]$.*

Proof.

Construction. Let $\mathcal{D} = \langle \Sigma, Q, \iota, \delta, \kappa \rangle$ be a DPA, where $\kappa : Q \to [1..k]$. We define the FDFA $\mathcal{F} = (\mathcal{Q}, \mathbf{P})$, where $\mathcal{Q}$ is the same as $\mathcal{D}$ (without acceptance), and each progress automaton $\mathcal{P}_q$ has k copies of $\mathcal{D}$, having $(q, \kappa(q))$ as its initial state, and moving to a j-th copy upon visiting a state with color j, provided that j is lower than the index of the current copy. The accepting states are those of the odd copies.

Formally: $\mathcal{Q} = \langle \Sigma, Q, \iota, \delta \rangle$, and for each state $q \in Q$, $\mathbf{P}$ has the DFA $\mathcal{P}_q = \langle \Sigma, Q \times [1..k], (q, \kappa(q)), \delta', F \rangle$, where for every $\sigma \in \Sigma$ and $i \in [1..k]$, $\delta'((q, i), \sigma) = (\delta(q, \sigma), \min(i, \kappa(\delta(q, \sigma))))$. The set F of accepting states is $\{Q \times \{i\} \mid i \text{ is odd }\}$.

Correctness. Analogous to the arguments in the proof of Theorem 5.3. □

From nondeterministic ω-automata. An NBA $\mathcal{A}$ can be translated into a saturated FDFA $\mathcal{F}$, by first determinizing $\mathcal{A}$ into an equivalent DPA $\mathcal{A}'$ [Pit06, FL15] (which might involve a $2^{O(n \log n)}$ state blowup and $O(n)$ colors [Sch09b]), and then polynomially translating $\mathcal{A}'$ into an equivalent FDFA (Theorem 5.4).

Proposition 5.5. *Let $\mathcal{B}$ be an* NBA *with n states. There is a saturated* FDFA *that characterizes $[\![\mathcal{B}]\!]$ with a leading automaton and progress automata of at most $2^{O(n \log n)}$ states each.*

A $2^{O(n \log n)}$ state blowup in this case is inevitable, based on the lower bound for complementing NBAs [Mic88, Yan06, Sch09a], the constant complementation of FDFAs, and the polynomial translation of a saturated FDFA to an NBA:

Theorem 5.6. *There exists a family of* NBA*s $\mathcal{B}_1, \mathcal{B}_2, \ldots$, such that for every $n \in \mathbb{N}$, $\mathcal{B}_n$ is of size n, while a saturated* FDFA *that characterizes $[\![\mathcal{B}_n]\!]$ must be of size (m, k), such that $\max(m, k) \geq 2^{\Omega(n \log n)}$.*

Proof. Michel [Mic88] has shown that there exists a family of languages $\{L_n\}_{n \geq 1}$, such that for every n, there exists an NBA of size n for L_n, but an NBA for L_n^c, the complement of L_n, must have at least $2^{n \log n}$ states.

Assume, towards a contradiction, that exist $n \in \mathbb{N}$ and a saturated FDFA $\mathcal{F}$ of size (m, k) that characterizes L_n, such that $\max(m, k) < 2^{\Omega(n \log n)}$. Then, by Theorem 4.1, there is a

saturated FDFA $\mathcal{F}^c$ of size (m, k) that characterizes L_n^c. Thus, by Theorem 5.8, we have an NBA of size smaller than $(2^{\Omega(n \log n)})^5 = 2^{\Omega(n \log n)}$ for L_n^c. Contradiction. □

5.2. **From FDFAs to ω-automata.** We show that saturated FDFAs can be polynomially translated into NBAs, yet translations of saturated FDFAs to DPAs may involve an inevitable exponential blowup.

To nondeterministic ω-automata. We show below that every saturated FDFA can be polynomially translated to an equivalent NBA. Since an NBA can be viewed as a special case of an NPA, a translation of saturated FDFAs to NPAs follows. Translating saturated FDFAs to NCAs is not always possible, as the latter are not expressive enough.

The translation goes along the lines of the construction given in [CNP94] for translating an $L_\$$-automaton into an equivalent NBA. We prove below that the construction is correct for saturated FDFAs, despite the fact that saturated FDFAs can be exponentially smaller than $L_\$$-automata.

We start with a lemma from [CNP94], which will serve us for one direction of the proof.

Lemma 5.7 ([CNP94])**.** *Let $M, N \subseteq \Sigma^*$ such that $M \cdot N^* = M$ and $N^+ = N$. Then for every ultimately periodic word $w \in \Sigma^\omega$ we have that $w \in M \cdot N^\omega$ iff there exist words $u \in M$ and $v \in N$ such that $uv^\omega = w$.*

We continue with the translation and its correctness.

Theorem 5.8. *For every saturated FDFA $\mathcal{F}$ of size (n, k), there exists an NBA $\mathcal{B}$ with $O(n^2k^3)$ states, such that $[\![\mathcal{F}]\!] \equiv [\![\mathcal{B}]\!]$.*

Proof.

Construction. Consider a saturated FDFA $\mathcal{F} = (\mathcal{Q}, \mathbf{P})$, where $\mathcal{Q} = \langle \Sigma, Q, \iota, \delta \rangle$, and for each state $q \in Q$, $\mathbf{P}$ has the progress DFA $\mathcal{P}_q = \langle \Sigma, P_q, \iota_q, \delta_q, F_q \rangle$.

For every $q \in Q$, let M_q be the language of finite words on which $\mathcal{Q}$ reaches q, namely $M_q = \{u \in \Sigma^* \mid \mathcal{Q}(u) = q\}$. For every $q \in Q$ and for every accepting state $f \in F_q$, let $N_{q,f}$ be the language of finite words on which $\mathcal{Q}$ makes a self-loop on q, $\mathcal{P}_q$ reaches f, and $\mathcal{P}_q$ makes a self-loop on f, namely $N_{q,f} = \{v \in \Sigma^* \mid (\delta(q, v) = q) \wedge (\mathcal{P}_q(v) = f) \wedge (\delta_q(f, v) = f)\}$. We define the ω-regular language

$$L = \bigcup_{\{(q,f) \mid (q \in Q) \wedge (f \in F_q)\}} M_q \cdot N_{q,f}^\omega \tag{5.1}$$

One can construct an NBA $\mathcal{B}$ that recognizes L and has up to $O(n^2k^3)$ states: L is the union of nk sublanguages; $\mathcal{B}$ will have nk corresponding subautomata, and will nondeterministically start in one of them. In each subautomaton, recognizing the language $M_q \cdot N_{q,f}^\omega$, a component of size n for M_q is obtained by a small modification to $\mathcal{Q}$, in which q can nondeterministically continue with an ϵ-transition[5] to a component realizing $N_{q,f}^\omega$. An NBA for the language $N_{q,f}$ consists of the intersection of three NBAs, for the languages $\{v \in \Sigma^* \mid \delta(q, v) = q\}$, $\{v \in \Sigma^* \mid \mathcal{P}_q(v) = f\}$, and $\{v \in \Sigma^* \mid \delta_q(f, v) = f\}$, each of which can be obtained by small modifications to either $\mathcal{Q}$ or $\mathcal{P}_q$, resulting in nk^2 states. Finally, the automaton for $N_{q,f}^\omega$ is obtained by adding ϵ-transitions in the automaton of $N_{q,f}$ from

[5]The ϵ-transitions can be removed from an NBA with no state blowup.

its accepting states to its initial state. Thus, each subautomaton is of size $n + nk^2$, and $\mathcal{B}$ is of size $nk(n + nk^2) \in O(n^2k^3)$.

Correctness. Consider an ultimately periodic word $uv^\omega \in [\![\mathcal{B}]\!]$. By the construction of $\mathcal{B}$, $uv^\omega \in L$, where L is defined by Equation (5.1). Hence, $uv^\omega \in M_q \cdot N_{q,f}^\omega$, for some $q \in Q$ and $f \in F_q$. By the definitions of M_q and $N_{q,f}$, we get that M_q and $N_{q,f}$ satisfy the hypothesis of Lemma 5.7, namely $N_{q,f}^+ = N_{q,f}$ and $M_q \cdot N_{q,f}^* = M_q$. Therefore, by Lemma 5.7, there exist finite words $u' \in M_q$ and $v' \in N_{q,f}$ such that $u'v'^\omega = uv^\omega$. From the definitions of M_q and $N_{q,f}$, it follows that the run of $\mathcal{Q}$ on u' ends in the state q, and $\mathcal{P}_q$ accepts v'. Furthermore, by the definition of $N_{q,f}$, we have $\delta(q, v') = q$, implying that (u', v') is the normalization of itself. Hence, $(u', v') \in [\![\mathcal{F}]\!]$. Since $\mathcal{F}$ is saturated and $u'v'^\omega = uv^\omega$, it follows that $(u, v) \in [\![\mathcal{F}]\!]$, as required.

As for the other direction, consider a pair $(u, v) \in [\![\mathcal{F}]\!]$, and let (x, y) be the normalization of (u, v) w.r.t. $\mathcal{Q}$. We will show that $xy^\omega \in L$, where L is defined by Equation (5.1), implying that $uv^\omega \in [\![\mathcal{B}]\!]$. Let $q = \mathcal{Q}(x)$, so we have that $\mathcal{P}_q(y)$ reaches some accepting state f of $\mathcal{P}_q$. Note, however, that it still does not guarantee that $y \in N_{q,f}$, since it might be that $\delta_q(f, y) \neq f$.

To prove that $xy^\omega \in L$, we will show that there is a pair $(x, y') \in \Sigma^{*+}$ and an accepting state $f' \in \mathcal{P}_q$, such that $y' = y^t$ for some positive integer t, and $y' \in N_{q,f'}$; namely $\delta(q, y') = q$, $\mathcal{P}_q(y') = f'$, and $\delta_q(f', y') = f'$. Note first that since $\mathcal{F}$ is saturated, it follows that for every positive integer i, $(x, y^i) \in [\![\mathcal{F}]\!]$, as $x(y^i)^\omega = xy^\omega$.

Now, for every positive integer i, $\mathcal{P}_q$ reaches some accepting state f_i when running on y^i. Since $\mathcal{P}_q$ has finitely many states, for a large enough i, $\mathcal{P}_q$ must reach the same accepting state $\hat{f}$ twice when running on y^i. Let h be the smallest positive integer such that $\mathcal{P}_q(y^h) = \hat{f}$, and r the smallest positive integer such that $\delta_q(\hat{f}, y^r) = \hat{f}$. Now, one can verify that choosing t to be an integer that is bigger than or equal to h and is divisible by r guarantees that $\delta(q, y^t) = q$ and $\delta_q(f', y^t) = f'$, where $f' = \mathcal{P}_q(y^t)$. □

To deterministic ω-automata. Deterministic Büchi and co-Büchi automata are not expressive enough for recognizing every ω-regular language. We thus consider the translation of saturated FDFAs to deterministic parity automata. A translation is possible by first polynomially translating the FDFA into an NBA (Theorem 5.8) and then determinizing the latter into a DPA (which might involve a $2^{O(n \log n)}$ state blowup [Mic88]).

Proposition 5.9. *Let $\mathcal{F}$ be a saturated* FDFA *of size (n, k). There exists a* DPA *$\mathcal{D}$ of size $2^{O(n^2k^3 \log n^2k^3)}$, such that $\mathcal{F} \equiv \mathcal{D}$.*

We show below that an exponential state blowup is inevitable.[6] The family of languages $\{L_n\}_{n\geq 1}$ below demonstrates the inherent gap between FDFAs and DPAs; an FDFA for L_n may only "remember" the smallest and biggest read numbers among $\{1, 2, ..., n\}$, using n^2 states, while a DPA for it must have at least 2^{n-1} states.

We define the family of languages $\{L_n\}_{n\geq 1}$ as follows. The alphabet of L_n is $\{1, 2, ..., n\}$, and a word belongs to it iff the following two conditions are met:

- A letter i is always followed by a letter j, such that $j \leq i + 1$. For example, $533245\ldots$ is a bad prefix, since 2 was followed by 4, while $55234122\ldots$ is a good prefix.

[6]This is also the case when translating FDFAs to deterministic Rabin [Rab69], Streett [Str82], and Muller [Mul63] automata, as explained in Remark 5.13.

- The number of letters that appear infinitely often is odd. For example, $2331(22343233)^\omega$ is in L_n, while $1(233)^\omega$ is not.

We show below how to construct, for every $n \geq 1$, a saturated FDFA of size polynomial in n that characterizes L_n. Intuitively, the leading automaton handles the safety condition of L_n, having $n+1$ states, and ensuring that a letter i is always followed by a letter j, such that $j \leq i+1$. The progress automata, which are identical, maintain the smallest and biggest number-letters that appeared, denoted by s and b, respectively. Since a number-letter i cannot be followed by a number-letter j, such that $j > i+1$, it follows that the total number of letters that appeared is equal to $b - s + 1$. Then, a state is accepting iff $b - s + 1$ is odd.

Lemma 5.10. *Let $n \geq 1$. There is a saturated* FDFA *of size $(n+1, n^2)$ characterizing L_n.*

Proof. We formally define an FDFA $\mathcal{F} = (\mathcal{Q}, \mathbf{P})$ for L_n over $\Sigma = \{1, 2, \ldots, n\}$, as follows.

The leading automaton is $\mathcal{Q} = (\Sigma, Q, \iota, \delta)$, where $Q = \{\bot, q_1, q_2, \ldots, q_n\}$; $\iota = q_n$; and for every $i, j \in [1..n], \delta(q_i, j) = q_j$ if $j \leq i+1$, and $\bot$ otherwise, and $\delta(\bot, j) = \bot$.

The progress automaton for the state $\bot$ consists of a single non-accepting state with a self-loop over all letters.

For every $i \in [1..n]$, the progress automaton for q_i is $\mathcal{P}_i = (\Sigma, P_i, \iota_i, \delta_i, F_i)$, where:

- $P_i = [1..n] \times [1..n]$
- $\iota_i = (n, 1)$
- δ_i: For every $\sigma \in \Sigma$ and $s, b \in [1..n], \delta_i((s,b), \sigma) = (\min(s, \sigma), \max(b, \sigma))$.
- $F_i = \{(s,b) \mid b - s \text{ is even } \}$

Notice that the progress automaton need not handle the safety requirement, as the leading automaton ensures it, due to the normalization in the acceptance criterion of an FDFA. □

A DPA for L_n cannot just remember the smallest and largest letters that were read, as these letters might not appear infinitely often. Furthermore, we prove below that the DPA must be of size exponential in n, by showing that its state space must be doubled when moving from L_n to L_{n+1}.

Lemma 5.11. *Every* DPA *that recognizes L_n must have at least 2^{n-1} states.*

Proof. The basic idea behind the proof is that the DPA cannot mix between 2 cycles of n different letters each. This is because a mixed cycle in a parity automaton is accepting/rejecting if its two sub-cycles are, while according to the definition of L_n, the mixed cycle should reject if both its sub-cycles accept, and vice versa. Hence, whenever adding a letter, the state space must be doubled.

In the formal proof below, we dub a reachable state from which the automaton can accept some word a *live state*. Consider a DPA $\mathcal{D}_n$ that recognizes L_n, and let q be some live state of $\mathcal{D}_n$. Observe that $[\![\mathcal{D}_n^q]\!]$, namely the language of the automaton that we get from $\mathcal{D}_n$ by changing the initial state to q, is the same as L_n except for having some restriction on the word prefixes. More formally, if a word $w \in [\![\mathcal{D}_n^q]\!]$ then $w \in L_n$, and if $w \in L_n$ then there is a finite word u, such that $uw \in [\![\mathcal{D}_n^q]\!]$. For every $n \in \mathbb{N}$, and every $u \in \Sigma^*$, let $L_{n,u} = \{w \mid uw \in L_n\}$ and let $\mathfrak{L}_n$ denote the set of languages $\{L_{n,u} \mid u \in \Sigma^*\}$. Note that there is actually only a finite number of prefixes u to consider (this follows e.g. from [MS97, Thm. 22]). Moreover, for every state q of $\mathcal{D}_n$ there is a corresponding word u_q such that $[\![\mathcal{D}_n^q]\!] = L_{n,u_q}$.

We prove by induction on n the following claim, from which the statement of the lemma immediately follows: Let $\mathcal{D}_n$ be a DPA over $\Sigma = \{1, 2, \ldots, n\}$ that recognizes some language in $\mathfrak{L}_n$. Then there are finite words $u, v \in \Sigma^*$, such that:

i) v contains all the letters in Σ;
ii) the run of $\mathcal{D}_n$ on u reaches some live state p; and
iii) the run of $\mathcal{D}_n$ on v from p returns to p, while visiting at least 2^{n-1} different states.

The base cases, for $n \in \{1, 2\}$, are trivial, as they mean a cycle of size at least 1 over v, for $n = 1$, and a cycle of size 2 for $n = 2$.

In the induction step, for $n \geq 2$, we consider a DPA $\mathcal{D}_{n+1}$ that recognizes some language $L \in \mathfrak{L}_{n+1}$. We shall define $\mathcal{D}'$ and $\mathcal{D}''$ to be the DPAs that result from $\mathcal{D}_{n+1}$ by removing all the transitions over the letter $n+1$ and by removing all the transitions over the letter 1, respectively.

Observe that for every state q that is live w.r.t. $\mathcal{D}_{n+1}$, we have that $[\![\mathcal{D}'^q]\!] \in \mathfrak{L}_n$, namely the language of the DPA that results from $\mathcal{D}_{n+1}$ by removing all the transitions over the letter $n+1$ and setting the initial state to q is in $\mathfrak{L}_n$. (Note that q might only be reachable via the letter $n+1$, yet it must have outgoing transitions over letters in $[2..n]$.) Analogously, $[\![\mathcal{D}''^q]\!]$ is isomorphic to a language in $\mathfrak{L}_n$ via the alphabet mapping of $i \mapsto (i-1)$. Hence, for every state q that is live w.r.t. $\mathcal{D}_{n+1}$, the induction hypothesis holds for $\mathcal{D}'^q$ and $\mathcal{D}''^q$.

We shall prove the induction claim by describing words $u, v \in \Sigma^*$, and showing that they satisfy the requirements above w.r.t. $\mathcal{D}_{n+1}$. We construct u by iteratively concatenating the words u'_i, v'_i, u''_i, and v''_i, which we define below, until the starting and ending states in some iteration k are the same. We then define the word v to be the last iteration, namely $u'_k\, v'_k\, u''_k\, v''_k$. Let q_1 be the initial state of $\mathcal{D}_{n+1}$. We define for every $i \in [1..k]$:

- u'_i and v'_i are the words that follow from the induction hypothesis on $\mathcal{D}'^{q_i}$, where q_i is the state that $\mathcal{D}_{n+1}$ reaches when reading $u'_1\, v'_1\, u''_1\, v''_1 \ldots u'_{i-1}\, v'_{i-1}\, u''_{i-1}\, v''_{i-1}$.
- u''_1 and v''_1 are the words that follow from the induction hypothesis on $\mathcal{D}''^{q'_i}$, where q'_i is the state that $\mathcal{D}_{n+1}$ reaches when reading $u'_1\, v'_1\, u''_1\, v''_1 \ldots u'_{i-1}\, v'_{i-1}\, u''_{i-1}\, v''_{i-1}\, u'_i\, v'_i$.

The word v obviously contains all the letters in Σ, as it is composed of subwords that contain all the letters in $\Sigma \setminus \{1\}$ and in $\Sigma \setminus \{n+1\}$. By the definition of u and v, we also have that the run of $\mathcal{D}_{n+1}$ on u reaches some live state p, and the run of $\mathcal{D}_{n+1}$ on v from p returns to p. Now, we need to prove that the run of $\mathcal{D}_{n+1}$ on v from p visits at least 2^n states.

We claim that when $\mathcal{D}_{n+1}$ runs on v from p, it visits disjoint set of states when reading v'_k and v''_k. This will provide the required result, as $\mathcal{D}_{n+1}$ visits at least 2^{n-1} states when reading each of v'_k and v''_k.

Assume, by way of contradiction, that $\mathcal{D}_{n+1}$ visits some state s both when reading v'_k and when reading v''_i. Let l' and r' be the parts of v'_k that $\mathcal{D}_{n+1}$ reads before and after reaching s, respectively, and l'' and r'' the analogous parts of v''_k. Now, define the infinite words $m' = u\, u'_k\, (l'\, r')^\omega$, $m'' = u\, u'_k\, l'\, (r''\, l'')^\omega$, and $m = u\, u'_k\, (l'\, r''\, l''\, r')^\omega$.

Observe that m' and m'' both belong or both do not belong to L, since there is the same number of letters (n) that appear infinitely often in each of them. The word m, on the other hand, belongs to L if m' and m'' do not belong to L, and vice versa, since $n+1$ letters appear infinitely often in it. However, the set of states that are visited infinitely often in the run of $\mathcal{D}_{n+1}$ on m is the union of the sets of states that appear infinitely often in the runs of $\mathcal{D}_{n+1}$ on m' and m''. Thus, if $\mathcal{D}_{n+1}$ accepts both m' and m'' it also accepts m, and if it rejects both m' and m'' it rejects m. (This follows from the fact that the minimal

number in a union of two sets is even/odd iff the minimum within both sets is even/odd.) Contradiction. □

Theorem 5.12. *There is a family of languages $\{L_n\}_{n\geq 1}$ over the alphabet $\{1, 2, \ldots, n\}$, such that for every $n \geq 1$, there is a saturated* FDFA *of size $(n + 1, n^2)$ that characterizes L_n, while a* DPA *for L_n must be of size at least 2^{n-1}.*

Proof. By Lemmas 5.10 and 5.11. □

Remark 5.13. A small adaptation to the proof of Lemma 5.11 shows an inevitable exponential blowup also when translating a saturated FDFA to a deterministic ω-automaton with a stronger acceptance condition of Rabin [Rab69] or Streett [Str82]: A mixed cycle in a Rabin automaton is rejecting if its two sub-cycles are, and a mixed cycle in a Streett automaton is accepting if its two sub-cycles are. Hence, the proof of Lemma 5.11 holds for both Rabin and Streett automata if proceeding in the induction step from an alphabet of size n to an alphabet of size $n + 2$, yielding a Rabin/Streett automaton of size at least $2^{\frac{n}{2}}$.

As for translating a saturated FDFA to a deterministic Muller automaton [Mul63], it is known that translating a DPA of size n into a deterministic Muller automaton might require the latter to have an accepting set of size exponential in n [Saf89]. (The family of languages in [Saf89] uses an alphabet of size exponential in the number of states of the DPA, however it can easily be changed to use an alphabet of linear size.) Hence, by Theorem 5.4, which shows a polynomial translation of DPAs to FDFAs, we get that translating an FDFA to a deterministic Muller automaton entails an accepting set of exponential size, in the worst case.

6. Discussion

The interest in FDFAs as a representation for ω-regular languages stems from the fact that they possess a correlation between the automaton states and the language right congruences, a property that traditional ω-automata lack. This property is beneficial in the context of learning, and indeed algorithms for learning ω-regular languages by means of saturated FDFAs were recently provided [AF14, LCZL16]. Analyzing the succinctness of saturated FDFAs and the complexity of their Boolean operations and decision problems, we believe that they provide an interesting formalism for representing ω-regular languages. Indeed, Boolean operations and decision problems can be performed in nondeterministic logarithmic space and their succinctness lies between deterministic and nondeterministic ω-automata.

References

[ABF16] D. Angluin, U. Boker, and D. Fisman. Families of DFAs as acceptors of ω-regular languages. In *41st International Symposium on Mathematical Foundations of Computer Science, MFCS 2016, August 22-26, 2016 - Kraków, Poland*, pages 11:1–11:14, 2016.

[AF14] D. Angluin and D. Fisman. Learning regular omega languages. In Peter Auer, Alexander Clark, Thomas Zeugmann, and Sandra Zilles, editors, *Algorithmic Learning Theory - 25th International Conference, ALT 2014, Bled, Slovenia, October 8-10, 2014. Proceedings*, volume 8776 of *Lecture Notes in Computer Science*, pages 125–139. Springer, 2014.

[AvMN05] Rajeev Alur, Pavol Černý, P. Madhusudan, and Wonhong Nam. Synthesis of interface specifications for Java classes. In *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '05, pages 98–109, New York, NY, USA, 2005. ACM.

[BHS+16] F. Blahoudek, M. Heizmann, S. Schewe, J. Strejcek, and M.H. Tsai. Complementing semi-deterministic büchi automata. In *Proc. of Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 9636 of *LNCS*, pages 770–787. Springer, 2016.

[Büc60] J.R. Büchi. Weak second-order arithmetic and finite automata. *Zeit. Math. Logik und Grundl. Math.*, 6:66–92, 1960.

[CCF+10] Yu-Fang Chen, Edmund M. Clarke, Azadeh Farzan, Ming-Hsien Tsai, Yih-Kuen Tsay, and Bow-Yaw Wang. Automated assume-guarantee reasoning through implicit learning. In *Computer Aided Verification, 22nd International Conference, CAV 2010, Edinburgh, UK, July 15-19, 2010. Proceedings*, pages 511–526, 2010.

[CCK+15] M. Chapman, H. Chockler, P. Kesseli, D. Kroening, O. Strichman, and M. Tautschnig. Learning the language of error. In *13th Int. Symp. on Automated Technology for Verification and Analysis*, pages 114–130, 2015.

[CNP94] H. Calbrix, M. Nivat, and A. Podelski. Ultimately periodic words of rational ω-languages. In *Proceedings of the 9th International Conference on Mathematical Foundations of Programming Semantics*, pages 554–566, London, UK, 1994. Springer-Verlag.

[CPPdR14] Georg Chalupar, Stefan Peherstorfer, Erik Poll, and Joeri de Ruiter. Automated reverse engineering using lego®. In *8th USENIX Workshop on Offensive Technologies (WOOT 14)*, San Diego, CA, August 2014. USENIX Association.

[FCTW08] A. Farzan, Y. Chenand E.M. Clarke, Y. Tsay, and B. Wang. Extending automated compositional verification to the full class of omega-regular languages. In C.R. Ramakrishnan and Jakob Rehof, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, volume 4963 of *Lecture Notes in Computer Science*, pages 2–17. Springer Berlin Heidelberg, 2008.

[FJV14] Paul Fiterau-Brostean, Ramon Janssen, and Frits W. Vaandrager. Learning fragments of the TCP network protocol. In *FMICS*, volume 8718 of *Lecture Notes in Computer Science*, pages 78–93. Springer, 2014.

[FKP11] Lu Feng, Marta Z. Kwiatkowska, and David Parker. Automated learning of probabilistic assumptions for compositional reasoning. In *Fundamental Approaches to Software Engineering - 14th International Conference, FASE 2011, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2011, Saarbrücken, Germany, March 26-April 3, 2011. Proceedings*, pages 2–17, 2011.

[FL15] D. Fisman and Y. Lustig. A modular approach for Büchi determinization. In Luca Aceto and David de Frutos-Escrig, editors, *26th International Conference on Concurrency Theory, CONCUR 2015, Madrid, Spain, September 1.4, 2015*, volume 42 of *LIPIcs*, pages 368–382. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.

[GJ79] M. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. MIT Press, 1979.

[Jon75] N.D. Jones. Space-bounded reducibility among combinatorial problems. *Journal of Computer and System Sciences*, 1975.

[Kla94] N. Klarlund. A homomorphism concept for omega-regularity. In L. Pacholski and J. Tiuryn, editors, *Computer Science Logic, 8th International Workshop, CSL '94, Kazimierz, Poland, September 25-30, 1994, Selected Papers*, volume 933 of *Lecture Notes in Computer Science*, pages 471–485. Springer, 1994.

[LCZL16] Yong Li, Yu-Fang Chen, Lijun Zhang, and Depeng Liu. A novel learning algorithm for Büchi automata based on family of DFAs and classification trees. *arXiv preprint arXiv:1610.07380*, 2016.

[McN66] R. McNaughton. Testing and generating infinite sequences by a finite automaton. *Information and Control*, 9(5):521–530, 1966.

[Mic88] M. Michel. Complementation is much more difficult with automata on infinite words. In *Manuscript, CNET*, 1988.

[MS97] O. Maler and L. Staiger. On syntactic congruences for omega-languages. *Theor. Comput. Sci.*, 183(1):93–112, 1997.

[Mul63] D.E. Muller. Infinite sequences and finite machines. In *Proc. 4th IEEE Symp. on Switching Circuit Theory and Logical design*, pages 3–16, 1963.

[Ner58] A. Nerode. Linear automaton transformations. *Proc. of the American Mathematical Society*, 9(4):541–544, 1858.

[NJ13] D. Neider and N. Jansen. Regular model checking using solver technologies and automata learning. In *NASA Formal Methods, 5th International Symposium, NFM 2013, Moffett Field, CA, USA, May 14-16, 2013. Proceedings*, pages 16–31, 2013.

[PGB+08] Corina S. Pasareanu, Dimitra Giannakopoulou, Mihaela Gheorghiu Bobaru, Jamieson M. Cobleigh, and Howard Barringer. Learning to divide and conquer: applying the l* algorithm to automate assume-guarantee reasoning. *Formal Methods in System Design*, 32(3):175–205, 2008.

[Pit06] N. Piterman. From nondeterministic Büchi and Streett automata to deterministic parity automata. In *21th IEEE Symposium on Logic in Computer Science (LICS 2006), 12-15 August 2006, Seattle, WA, USA, Proceedings*, pages 255–264. IEEE Computer Society, 2006.

[PVY02] D. Peled, M. Vardi, and M. Yannakakis. Black box checking. *Journal of Automata, Languages and Combinatorics*, 7(2):225–246, 2002.

[Rab69] M.O. Rabin. Decidability of second order theories and automata on infinite trees. *Transaction of the AMS*, 141:1–35, 1969.

[RMSM09] Harald Raffelt, Maik Merten, Bernhard Steffen, and Tiziana Margaria. Dynamic testing via automata learning. *STTT*, 11(4):307–324, 2009.

[Saë90] B. L. Saëc. Saturating right congruences. *ITA*, 24:545–560, 1990.

[Saf89] S. Safra. *Complexity of automata on infinite objects*. PhD thesis, Weizmann Institute of Science, 1989.

[Sch09a] S. Schewe. Büchi complementation made tight. In *Proc. 26th Symp. on Theoretical Aspects of Computer Science (STACS)*, volume 3 of *LIPIcs*, pages 661–672, 2009.

[Sch09b] S. Schewe. Tighter bounds for the determinization of Büchi automata. In *Proc. 12th Int. Conf. on Foundations of Software Science and Computation Structures (FoSSaCS)*, pages 167–181, 2009.

[Str82] R.S. Streett. Propositional dynamic logic of looping and converse. *Information and Control*, 54:121–141, 1982.

[VSVA05] Abhay Vardhan, Koushik Sen, Mahesh Viswanathan, and Gul Agha. Using language inference to verify omega-regular properties. In *Tools and Algorithms for the Construction and Analysis of Systems, 11th International Conference, TACAS 2005, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2005, Edinburgh, UK, April 4-8, 2005, Proceedings*, pages 45–60, 2005.

[VW86] M. Y. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification (preliminary report). In *LICS*, pages 332–344. IEEE Computer Society, 1986.

[Wil91] T. Wilke. *An Eilenberg theorem for ∞-languages*, pages 588–599. Springer Berlin Heidelberg, Berlin, Heidelberg, 1991.

[Wil93] T. Wilke. An algebraic theory for regular languages of finite and infinite words. *IJAC*, 3(4):447–490, 1993.

[WN05] Westley Weimer and George C. Necula. Mining temporal specifications for error detection. In *Tools and Algorithms for the Construction and Analysis of Systems, 11th International Conference, TACAS 2005, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2005, Edinburgh, UK, April 4-8, 2005, Proceedings*, pages 461–476, 2005.

[YAL14] Yifei Yuan, Rajeev Alur, and Boon Thau Loo. Netegg: Programming network policies by examples. In *Proceedings of the 13th ACM Workshop on Hot Topics in Networks, HotNets-XIII, Los Angeles, CA, USA, October 27-28, 2014*, pages 20:1–20:7, 2014.

[Yan06] Q. Yan. Lower bounds for complementation of ω-automata via the full automata technique. In *Proc. 33rd Int. Colloq. on Automata, Languages, and Programming (ICALP)*, volume 4052 of *LNCS*, pages 589–600, 2006.

This work is licensed under the Creative Commons Attribution-NoDerivs License. To view a copy of this license, visit http://creativecommons.org/licenses/by-nd/2.0/ or send a letter to Creative Commons, 171 Second St, Suite 300, San Francisco, CA 94105, USA, or Eisenacher Strasse 2, 10777 Berlin, Germany